

UNCLASSIFIED

AD 89515

Armed Services Technical Information Agency

Reproduced by

DOCUMENT SERVICE CENTER

KNOTT BUILDING, DAYTON, 2, OHIO


This document is the property of the United States Government. It is furnished for the duration of the contract and shall be returned when no longer required, or upon recall by ASTIA to the following address: Armed Services Technical Information Agency, Document Service Center, Knott Building, Dayton 2, Ohio.

NOTICE: WHEN GOVERNMENT OR OTHER DRAWINGS, SPECIFICATIONS OR OTHER DATA ARE USED FOR ANY PURPOSE OTHER THAN IN CONNECTION WITH A DEFINITELY RELATED GOVERNMENT PROCUREMENT OPERATION, THE U. S. GOVERNMENT THEREBY INCURS NO RESPONSIBILITY, NOR ANY OBLIGATION WHATSOEVER; AND THE FACT THAT THE GOVERNMENT MAY HAVE FORMULATED, FURNISHED, OR IN ANY WAY SUPPLIED THE SAID DRAWINGS, SPECIFICATIONS, OR OTHER DATA IS NOT TO BE REGARDED BY IMPLICATION OR OTHERWISE AS IN ANY MANNER LICENSING THE HOLDER OR ANY OTHER PERSON OR CORPORATION, OR CONVEYING ANY RIGHTS OR PERMISSION TO MANUFACTURE, USE OR SELL ANY PATENTED INVENTION THAT MAY IN ANY WAY BE RELATED THERETO.

UNCLASSIFIED


#30
FC

AD No. 89 5/3
ASTIA FILE COPY



Carnegie Institute
of Technology

Pittsburgh 13, Pennsylvania



GRADUATE SCHOOL of INDUSTRIAL ADMINISTRATION

William Larimer Mellon, Founder

A LINEAR DECISION RULE FOR
PRODUCTION AND EMPLOYMENT SCHEDULING

by

Charles C. Holt, Franco Modigliani
and Herbert A. Simon*

Approved for Release by
the President John F. Kennedy
any person who communicates
information to the President
76001, Project 101-01-001.

May 1955

Graduate School of Industrial Administration
Carnegie Institute of Technology

*Research undertaken for the project Planning and Control of Industrial Operations, under a grant from the Office of Naval Research. This paper was written by Mr. Holt, but it rests on research done by all three authors. The authors wish to express their appreciation for the effective assistance of their colleague, Robert W. Culbertson, and of the following graduate students: John F. Muth, Clyde E. Roberson, Robert F. Byrne, Toshiro Makibuchi, David V. Heebink, William P. Maughan, and Lawrence Schwartz.

Special thanks are due the paint company whose full cooperation with our research made this paper possible.

Table of Contents

	<u>Page Numbers</u>
I. INTRODUCTION	1
II. THE DECISION PROBLEM:	
SCHEDULING PRODUCTION RATE AND WORK FORCE	3
The Costs of Responding to Fluctuations in Orders	5
Errors in Forecasting Orders	8
The Time Sequence of Decisions	9
III. A QUADRATIC CRITERION FOR THE SCHEDULING DECISIONS	
IN A PAINT FACTORY	10
Regular Payroll, Hiring, and Layoff Costs	12
Overtime Costs	14
Inventory, Back Order and Machine Setup Costs	17
The Cost Function for the Paint Factory	20
IV. THE OPTIMAL DECISION RULES FOR THE PAINT FACTORY	24
V. COMPARISON OF DECISION PERFORMANCES BY THE FACTORY	
AND THE DECISION RULES	29
History of Factory Operations under Alternative	
Decision Methods	31
Cost Comparisons under Alternative Decision Methods	34
VI. CONCLUSION	40
<u>APPENDIX:</u> Derivation of the Conditions for Minimum Costs	43

I Introduction

The decision problems involved in setting the aggregate production rate of a factory and setting the size of its work force are frequently both complex and difficult. The quality of these decisions can be of great importance to the profitability of an individual company, and when viewed on a national scale, these decisions have a significant impact on the efficiency of the economy as a whole. This paper reports some of the findings of a research team that has been developing new methods to enable production executives to make better decisions and to make them more easily than they can with prevailing procedures. With the cooperation of a manufacturing concern, the new methods have been developed in the context of a set of concrete production scheduling problems that were found in a factory operated by the company.

The new method, published for the first time in this paper,* involves: (1) formalizing and quantifying the decision problem (using a quadratic criterion function) and (2), calculating a generalized optimal solution of the problem in the form of a (linear) decision rule. Like a rule of thumb, an optimal decision rule prescribes a course of action when it is applied to a particular set of circumstances; but, unlike most rules of thumb, an optimal decision rule prescribes courses of action for which the claim can

* Research that served as background for the present work was reported in "Optimal Decision Rules for Production and Inventory Control," by C. C. Holt and H. A. Simon, in Proceedings of the Conference on Operations Research in Production and Inventory Control, January 1954, Case Institute of Technology.

be made that the decisions are "the best possible," and the meaning of "best" is clearly specified. The ultimate test, of course, must be whether the new decision methods do or do not outperform prevailing decision methods when full allowance is made for the cost of obtaining the optimal rules.

In the body of this paper, we explore the problem of setting the aggregate production rate and size of the work force. We describe the particular form that this problem takes in the factory of the cooperating company, including a consideration of the various types of costs and intangible penalties that are relevant in making the decision. Then, without going into details about the methods used to solve the problem (these are contained in a technical appendix), we present the solution in the form of the decision rule that is optimal for the type of decision criterion that was used. We found that, once the decision problem of the cooperating company was formalized and quantified, the numerical constants appearing in the decision rule could be computed with a desk calculator in three man-hours.

After this decision rule was obtained, it was then applied to the monthly production rate and labor force decisions that faced the company over a six-year period. Using the decision rule, each of these monthly decisions required only a five-minute calculation. Comparisons are presented of the actual performance of the factory with the hypothetical performance — the performance that would have been realized if the new methods had been used. These performances are compared also by means of cost estimates.

The decision method that is here applied to a particular factory, should be directly applicable to other factories having the same kinds of costs,

The method presented in this paper may^{also} be adapted readily to factories with types of costs entirely different from those in the example presented. However, until the techniques for applying the method have been further developed, each new application will undoubtedly require some developmental effort. Ultimately, decision criteria that can be adequately approximated by quadratic functions, and linear decision rules should prove useful in handling a wide range of decision-making problems quite beyond the specific problem of production scheduling.

II The Decision Problem: Scheduling Production Rate and Work Force

It is important at the outset to outline clearly the many facets of the decision problem that faces an executive in setting the aggregate production rate and size of the labor force of a factory. A good place to start is to define the variables whose scheduling constitutes the decision problem at hand. By aggregate production rate we mean a physical measure of production per unit of time (per week or per month, for example). Most factories produce many products, rather than just one; hence, a common unit must be found for adding quantities of different products. For example, a unit of weight, volume, work required, or value might serve as a suitable common denominator. The other decision variable, work force, refers to the number of employees to whom there is a company commitment to supply regular work.

The initial limitation of the problem, to consider these two decision variables only, requires comment. Clearly neither decision can be separated

completely from other decisions about product mix, labor mix, and production sequences. For example, the number of workers needed may depend on the number of different products to be produced as well as the aggregate production rate. Although our limitation of the decision problem rules out of consideration certain interactions that will be important for some factories, this limitation appeared reasonable in order to keep the initial research problem within reasonable bounds. In applying the decision rule to the cooperating factory, auxiliary techniques, which will not be described here, have been employed.^{1/}

The problem is to choose a course of action that will produce the results that are desired. Choices are "problematic" when complexities keep the best course of action from being obvious. In deciding upon the production rate and the labor force of a factory there are three important aspects that contribute sufficient complexity to constitute a formidable problem: 1) How should production and employment be adjusted to fluctuations in the orders received? 2) What provision should be made for errors in the forecasts of future orders? and 3) What is the implication of the fact that the current decision is but one of a sequence of decisions to be made at successive points of time? We consider each of these questions in turn.

^{1/} See "Some Techniques for the Solution of Dynamic Programming Problems in Production Scheduling," by Herbert A. Simon, Charles C. Holt, and Franco Modigliani, to be published in the Journal of the Society for Industrial and Applied Mathematics. (ONR Memorandum No. 29)

The Costs of Responding to Fluctuations in Orders

If the customers of a factory placed their orders in such a way as to call for a constant flow of shipments of finished product, the two decisions under consideration would hardly constitute a problem. In actual fact, orders (or more precisely, ordered shipments) are subject to substantial fluctuation, and the question arises as to how these fluctuations should be "absorbed." That the problem is not trivial may be seen by considering three "pure" alternative ways of responding to such fluctuations.

These alternatives are: (1) to adjust the size of the work force by hiring and firing in exact conformity with the fluctuations in orders; (2) to adjust the production rate into conformity with orders by working overtime or "undertime" with a constant work force; and (3) to allow inventory and the backlog of orders to fluctuate while maintaining a constant work force and a constant production rate. Each of these "pure" alternatives has certain costs -- interpreting that term broadly to include any tangible or intangible penalty -- associated with it.

(1) Under the first alternative, an increase in orders would be met by hiring, while a decrease in orders would be accompanied by layoffs. While this procedure is clearly not optimal for the economy as a whole, since the numbers of workers is constant in the short run, it is nevertheless an admissible alternative for an individual company. However, training and reorganization are usually required when the work force is expanded;

and terminal pay, bumping,^{1/} and loss of worker morale frequently occur when the labor force is contracted. Since plant and equipment are fixed in the short run, increases in the work force may decrease labor productivity. This cost can be avoided by maintaining the plant and equipment necessary for peak employment, or by paying the premiums involved in second and third shift operation. A similar problem of unbalance may arise when the total work force fluctuates, but some components of the work force, supervision for example, cannot easily be changed. For all these reasons, fluctuations in the work force are costly. From work force considerations alone, the "ideal" work force would be one of constant size, with an optimum balance of men, machines and supervision.

(2) The second alternative would realize this "ideal" work force situation by absorbing fluctuations in orders with corresponding fluctuations in overtime work without changing the size of the work force. However, since there is an upper limit to what a worker can produce by working overtime, the necessity for meeting peak orders would govern the size of the work force. When orders fall to lower levels overtime is eliminated, but with a further fall in orders "undertime" occurs, i.e., there is not enough productive work to keep the work force busy throughout the regular work week. Hence, this alternative has its limitations. The well-recognized costs of the overtime premium do not require emphasis; the cost of "undertime"

^{1/} Union seniority rules sometimes require a whole sequence of job transfers when a single job is eliminated.

is less obvious. Man-hours paid for with no product output constitute a cost to the factory unless "fill-in" jobs (e.g., maintenance) can be scheduled, or leisure on the job has an important positive morale value. Sometimes the cost of undertime can be passed on to the employees by shortening the work week, but even here it is unlikely that the company completely escapes indirect penalties. Thus absorbing fluctuations in orders by overtime and undertime incurs various penalties and costs. From overtime and undertime considerations alone neither "should" be incurred; "ideal" overtime and "ideal" undertime are zero.

(3) Finally, the fluctuations in orders may be absorbed by allowing the inventory of finished goods to fluctuate or, lacking a finished inventory, by allowing the backlog of unfilled orders to fluctuate. Big upward swings in inventory necessitate large storage facilities, large amounts of working capital and other direct costs, and create risks such as obsolescence. Big downward swings of inventory, culminating in large order backlogs, impose intangible costs on the company — poor delivery service to customers/^{may} lead to loss of sales. Clearly, absorption of order fluctuations by building up or drawing upon inventory (considering an order backlog as a negative inventory) is not altogether a happy answer. If only inventory costs are taken into account the output of the factory should exactly match the shipments to be made; finished inventory "should" be zero!

It is abundantly clear that the fluctuations in customers' orders impose costs and penalties on the supplying company regardless of which

policy alternative it may follow in responding to these fluctuations. Because orders fluctuate, these dynamic costs are relevant and important in production and labor force decisions. Or stated differently, when a factory must absorb fluctuations in shipments imposed by its customers' orders, every alternative for absorbing these fluctuations has associated with it a set of costs and penalties for the company. In order to make a good decision, these costs must be weighed to determine what kind of policy will minimize them.

In general, none of the pure alternatives discussed above will prove best, but rather some carefully weighted combination of them. Order fluctuations should, in general, be absorbed partly by inventory, partly by overtime, and partly by hiring and layoffs, and the best allocation among these parts will depend upon the costs in each particular factory. But even for a particular factory, the best allocation is not fixed, but will vary with the severity (frequency) of the fluctuations.

Despite the fact that countless production executives are faced daily with this allocation problem, very little work has been done to find the optimal policy even for the case where fluctuations in orders are highly predictable, as with seasonal fluctuations. Unfortunately, however, the problem is even more difficult, for fluctuations of orders can seldom be foreseen accurately. This brings us to the problem of forecast errors.

Errors in Forecasting Orders

Any decision setting the production rate and work force of a factory will appear in retrospect to have been a good or poor decision depending

upon what orders were in fact received after the decision was made. A decision is not good or bad in itself, but only relative to the state of the world during the time in which the influence of the decision is being felt. Of course, the future state of the world -- in our case future receipts of orders -- ordinarily cannot be known in advance exactly. Consequently the decision must be made in a setting of uncertainty. At the time a decision has to be made, the outcomes associated with each of the alternatives are uncertain, since they depend partly on the unknown future. The better the future can be forecasted, the less uncertainty is involved in a decision; but uncertainty inevitably enters ^{the} decision to some extent, and must be resolved in one way or another.

It is useful to distinguish two aspects of the forecasting problem:

- 1) With a given forecast, produced by methods whose accuracy in the past is known, how should the decision be reached (i.e., how should decisions be affected by the fact that the forecasts are known to be subject to error)?
- 2) For any given forecast method, how large are the costs incurred as the direct result of its forecast errors? Knowledge of forecast accuracy usually is important both in using the forecasts and in selecting the forecasting method. However, the most accurate forecast method is not always the best, since the cost of obtaining the forecasts may exceed their value in improving the quality of decisions.

The Time Sequence of Decisions

The decisions setting production rate and work force fortunately do not involve a once-and-for-all commitment, but rather permit successive

review and revision as the passage of time provides new information. The errors of past forecasts are observed, and new information is obtained that provides a basis for revised forecasts. A decision once taken commits the production executive only until a new decision is made.^{1/} Although a decision based on an erroneous forecast can to a large extent be offset by subsequent decisions, such oscillations incur the same types of costs as do fluctuations in orders. No one decision is good or bad in itself, but only in its relation to the preceding and following decisions, and the preceding and following orders. Thus it is clear that the time sequence of decisions is an important aspect of the scheduling problem.

Having outlined in a rough way the major components of an important decision problem, and having indicated that the decisions depend upon interacting factors sufficiently complex to make the choice extremely difficult, we next consider a practicable method for finding a solution to this problem.

III A Quadratic Criterion for the Scheduling Decisions in a Paint Factory

Rather than present the new method in its most general form, we will describe an actual case that we have studied in detail, namely, the paint factory whose scheduling problems supplied stimulus to the development of the method. Aside from the value of a concrete example, we believe that

^{1/} The selection of the optimal decision period is not studied in this paper, but it is known that the optimal period depends on the size of forecast errors, the cost of forecasting, the time required for enough new information to accumulate to improve forecasts previously made, the cost of making and administering new decisions, and the relative cost of making many small decision changes versus making a smaller number of larger changes.

this case is representative of the scheduling problems in a large number of factories to which the same method, even in its details, may be applicable.

A decision-making problem of a business firm may usually be stated formally as a problem of finding a maximum (or minimum) of some criterion. Sometimes profit is the criterion to be maximized; in most cases profit will at least have considerable weight in the criterion function. In the paint factory, we treat the scheduling of production and employment from the point of view of the production manager. We assume that sales volume and price are beyond his control, so that revenue is a given, and the maximization of profits becomes the minimization of costs. We should emphasize that "costs" are interpreted broadly to include any relevant considerations to which the decision maker chooses to attach importance.

In order to apply the method, all costs, even though some are intangibles, must be reduced to quantitative terms and expressed in comparable units — presumably dollars. We can sometimes attach a dollar value to intangible factors by asking how much the management would be willing to spend outright in order to change these factors. To be sure, difficulties arise in quantifying a criterion function; but no system of rational decision-making can escape the task of assigning weights to the objectives that are desired.

In order to translate the scheduling problem into a mathematical problem of minimizing a cost function, we need a mathematical form that is both sufficiently flexible to approximate a wide range of complex cost

relationships, and sufficiently simple to allow easy mathematical solution. From consideration of the kinds of costs that are involved in the scheduling problem it appears that a U-shaped cost curve is required. For example, the cost of inventory is high when inventory is large, and high also at the other extreme when inventory is so small that there are frequent run-outs of particular products which cause back orders and a high penalty for delayed shipments to customers. Somewhere between these extremes, the combined costs are at a minimum. With these considerations in view, we decided that the cost function could reasonably be approximated by a sum of linear and squared terms in the controlled and uncontrolled variables -- technically, by a positive definite quadratic form.

In the following pages we will analyze the costs that are important in the particular paint factory that has been studied, and then show that these can be approximated by a quadratic cost function. Decisions are assumed to be made at regular time intervals (in this case monthly), rather than continuously or intermittently, and the costs are expressed as costs per month. It is convenient to relate these costs to the three alternative ways, discussed earlier, of absorbing order fluctuations.

It should be emphasized again that the following application of the new method represents a special case -- the method itself is far more general.

Regular Payroll, Hiring, and Layoff Costs

When order fluctuations are absorbed by increasing and decreasing the work force, the following costs are affected: regular payroll, hiring, and layoff costs.

The size of the work force is adjusted once a month, and setting the work force at a certain level implies a commitment to pay these employees their regular time (as contrasted with overtime) wages for a month. This is shown in Figure 1 by the solid line which may be represented algebraically by the linear cost function, Equation 1.. (In the equations that follow, the C's represent constants.)


The other labor costs mentioned are associated not with the size of the work force, but with changes in its size. The cost of hiring and training people rises with the number hired, as indicated by the solid line plotted in Figure 2. The cost of laying off workers derives from terminal pay, reorganization, etc., and rises with the number of workers laid off. The cost incurred each month depends on the change in the size of the work force between successive months. Since these costs increase both with increases and decreases in the work force, the quadratic curve represented by Equation 2 is a suitable first approximation.

Random factors may affect the costs of hiring and firing, e.g., how much difficulty is experienced in a particular case in hiring a man of desired qualifications, or how much reorganization is required in making a particular reduction in work force. Consequently the cost curve should be viewed as a curve of the average (expected) cost of changes of various sizes in the work force.

Whether these costs actually rise at an increasing or decreasing rate is difficult to determine. It can be argued that reorganization costs are more than proportionately larger for large layoffs than for small

Figure 1

REGULAR PAYROLL COST




APPROXIMATING COST FUNCTION

Eq. 1) Regular Payroll Cost = $C_1 W$

Figure 2

HIRING AND LAYOFF COSTS



APPROXIMATING COST FUNCTION — — —

Eq. 2) Cost of Hiring and Layoffs = $C_2(W_t - W_{t-1})^2$

layoffs; and similarly the efficiency of hiring, measured in terms of the quality of the employees hired, may fall when a large number of people are hired at one time. If this argument holds, then the quadratic curve is especially suitable. But if not, the quadratic still can give a tolerable approximation over a range. The parameters of the function should be set at the values that will give the best possible approximation to the cost curve over the range in which changes in the work force are expected to fluctuate.^{1/}

Overtime Costs

If order fluctuations are absorbed by increasing and decreasing production without changing the work force, then overtime and undertime costs are incurred. Overtime involves wage payments at an hourly rate fifty per cent higher than is paid for regular time. Undertime is a waste of labor time that is paid for in the regular payroll, but is not used for productive activities.^{2/} However, unlike overtime, undertime does not incur an increase in out-of-pocket expense.

The cost of overtime depends on two decision variables, the size of the work force, W , and the aggregate production rate, P . The simplest form

^{1/} Note the implied circularity. In order to obtain optimal decisions we need initially to know optimal fluctuation amplitudes of controlled variables. But for practical purposes we need to know only the general range of fluctuations, which can be estimated to a sufficiently close approximation.

^{2/} It may be possible to perform maintenance activities with labor that would otherwise be wasted. If so, this possibility should be taken into account.

of this cost relation is shown in Figure 3. With a given work force, W_1 , and an average worker productivity, K , the expression $K W_1$, is the maximum number of units that can be produced in a month without incurring any overtime. In order to produce at higher rates than $K W_1$, overtime is required, and its amount increases with increased production.

The relationship shown in Figure 3 can be expected to occur only if there are no discontinuities and no random disturbances in the production process. However, these are usually present, and should be taken into account. For example, since workers are each somewhat specialized in function, it is likely that a small increase in production would require only a few employees who work in bottleneck functions to work overtime. As production is increased further, more and more employees are required to work overtime until the whole work force is doing some overtime work. The effect of this is to smooth the overtime cost curve of Figure 3 to that shown in Figure 4.

Random disturbances have the same effect of smoothing the overtime curve. For example, given the number of units to be produced in a month, the total number of man-hours that will be required is not uniquely determined in advance, but will be affected by numerous random disturbances, such as machine breakdowns, quality control problems, productivity fluctuations, etc. Overtime is determined by the excess of the hours that prove to be required by the production target over and above the number of regular-time hours available from the work force in the month. Since the production and employment schedule is made before there is knowledge of the particular

Figure 3

OVERTIME COST WHEN THE WORK FORCE IS OF SIZE, W_1





Figure 4

OVERTIME COST WHEN THE WORK FORCE IS OF SIZE, W_1
AND RANDOMNESS AND DISCONTINUITIES ARE PRESENT



APPROXIMATING COST FUNCTION — — — —

Eq. 4) Expected Cost of Overtime = $C_3(P - C_4 W_1)^2 + C_5 P - C_6 W_1$

disturbances that will occur during the month, estimated overtime costs must depend on an estimate of the probabilities that such disturbances will occur. This probability distribution smooths the curve of expected overtime cost shown in Figure 4. The higher the production target with a given size work force, the greater is the probability that some disturbance will occur that will necessitate overtime work to get out the specified production.

In setting the production rate and the work force for a month, it is not certain in advance whether overtime or undertime will occur. In order for the scheduling decision to minimize costs, the cost of having a larger work force than might prove to be needed must be weighed against the cost of having a smaller and cheaper work force, but then perhaps finding it necessary to pay for considerable overtime.

The quadratic curve that approximates the expected cost of overtime for a given size, W_1 , of work force, and for different production rates is shown by Equation 4. As production, P , exceeds $C_4 W_1$, a level set by the size of the work force, overtime costs increase. The linear terms, $C_5 P$ and $C_6 W$, are added to improve the approximation.

The foregoing discussion was premised on a given work force, W_1 , but clearly the size of the work force can change. Hence there is a whole family of cost curves similar to that shown in Figure 4, one for each size of work force. This family of overtime cost curves is obtained by substituting other values for W_1 in Equation 4.

Inventory, Back Order and Machine Setup Costs

When order fluctuations are absorbed by inventory and back-order fluctuations other costs are incurred. Increased inventory increases the costs of interest, obsolescence, handling, storage, etc. The decrease of inventory to avoid these costs increases the probability of running out of individual products, thereby incurring the penalty of delaying customer shipments and possibly losing sales. Also, as aggregate inventory is reduced, the average production batch size should be decreased in order to maintain a balanced inventory; consequently, the cost of additional machine setups is incurred. An analysis of the total of these costs will indicate the optimal level of aggregate inventory at which these costs are minimum.

Production decisions in the paint factory are to be made monthly, and prior to each decision the aggregate inventory position should be observed. In formulating the cost function, we assume that the inventory and back order position at the end of each month is representative of the average inventory and back-order positions during the month, and consequently may reasonably be used to estimate the costs related to inventory that were incurred during the month. If this assumption is not tenable, it probably indicates that production decisions should be made more frequently than once a month. Production that is scheduled for a month is assumed to be completed during the month.^{1/}

^{1/} Production processes requiring several decision periods to complete may be accommodated in the mathematical model, but this was not necessary in the paint factory.

In order to have a simple relation between a month's production and the inventory at the end of the month, it is convenient to use the variable, net inventory, defined as inventory minus back orders. Net inventory is increased by production, regardless of whether the paint is added to physical inventory or shipped out to decrease the number of back orders. The paint factory usually ships immediately upon receipt of an order, and orders not so shipped are treated as back orders. Consequently, net inventory is affected immediately upon receipt of an order.^{1/}

Familiar lot size formulas^{2/} may be used to determine the optimal production batch size for each paint and the optimal safety stock to protect against its running out while a new batch is being produced. These formulae rest on plausible assumptions about the costs of holding inventory, the cost of back orders, and the probability of errors in forecasting orders for the particular paint. By adding, for each paint, the optimal average safety stock to one-half the optimal batch size we obtain its optimal average inventory.

Then by adding together these optimal average inventories for all the paints that are stocked, we obtain an optimal aggregate inventory for the

1/ For many factories a lead time is allowed between the receipt of an order and the shipping date requested by the customer. In such a case an order would not affect net inventory until the ordered shipping date. However, the receipt of an order supplies vital information by enabling a perfect forecast to be made of future shipments over a lead time horizon.

2/ See T. M. Whitin, The Theory of Inventory Management; and K. J. Arrow, T. Harris, J. Marschak, "Optimal Inventory Policy," Econometrica, Vol. 19, No. 3, July 1951.


whole factory. To convert this optimal aggregate inventory to the corresponding optimal net inventory, we need to subtract the total back orders for all paints that would be expected to occur, on the average, when the inventory is at its optimal level.

From lot size formulas it is known that both the optimal batch size and the optimal safety stock increase roughly with the square root of the order rate of the individual paint. Thus the optimal aggregate inventory must increase with increased aggregate order rate (total shipments ordered per month). The total expected back orders corresponding to any given size of inventory must also increase with an increased aggregate order rate. By combining these two relationships it appears that optimal net inventory increases with the aggregate order rate. The relationship between optimal net inventory and aggregate order rate may be approximated^{1/} over a range by a function of the form: optimal net inventory = $C_8 + C_9 O$ where the C 's are constants, and O is the aggregate order rate.

When actual net inventory deviates from the optimal net inventory, $(C_8 + C_9 O)$, in either direction, costs rise as shown in Figure 5. If net inventory falls below this optimal level, then the safety stocks and batch sizes of individual paints must be reduced. We assume that these

1/ Since back orders will generally be small relative to inventory, the square root relation between aggregate inventory and order rate dominates the relationship between net inventory and order rate. Over a limited range a square root function can be approximated by a linear one.

Figure 5



APPROXIMATING COST FUNCTION — — —

Eq. 5) Expected Inventory, Back Order, and Setup Costs

$$= C_7 [I - (C_8 + C_9 O)]^2$$

I = Net Inventory = Inventory - Back Orders

O = Order Rate (Units of Product whose Shipment was Ordered during the Month)

reductions are optimally distributed over the individual paints by some procedure for scheduling the production of individual products within the constraint of the aggregate production decision.^{1/} The rise in costs as net inventory declines can be estimated by costing the increased number of machine setups, the increased back orders and the decreased inventory. A similar cost calculation can be made for the situation in which net inventory is above the optimal level. In this way the relation, which is shown by the solid line in Figure 5, between expected costs and net inventory, can be determined. Over a range, the curve of inventory-related costs may be approximated adequately by a quadratic of the form shown in Equation 5 in which cost rises as the square of the deviation of net inventory from the optimal level, $(C_8 + C_9 O)$.

The Cost Function for the Paint Factory

Having examined the individual cost components we can now construct the complete cost function for production and employment scheduling. Since the objective is to schedule production and employment in such a way as to minimize costs, we need a cost function that adds together all the component costs that have been discussed above. Since each month's decision has cost implications that extend over an appreciable length of time, this cost function must span a sufficient time to include virtually all of the cost implications of the decision. The first requirement is met by adding all of the costs attributable to each month; and the second, by adding all of

^{1/} Methods for scheduling the production of individual products is under study, and we will report on this work in a later paper.

these monthly costs over an extended period of time. The discounting of costs that occur at different points of time by means of an interest rate factor is neglected as an unessential refinement.

Since future costs depend on future sales they are, of course, uncertain. This problem is met by calculating what the costs would be for each combination of forecast errors and taking a weighted average of these costs, using probability weights. This expected cost is to be minimized. No consideration is given to the variability of costs, but only their long term total. The decision problem can now be stated formally.

The optimal production and employment decisions are those that minimize the expected value of total cost, C_N . (see Equation 6). This cost is the sum of the costs attributable to N months as shown in Equation 7. The total cost attributable to one month, C_t , is shown in Equation 8 to be the sum of the component costs that have previously been discussed. Note that the time subscript, t , has been added to indicate that the variables may take on different values at different points of time. Equation 9 shows the relationship between inventory at the beginning of the month, production during the month, sales during the month, and the month-end inventory. This relationship, of course, applies to each month and must be taken into account in minimizing costs.

Eq. 6) Find the decisions that minimize $E(C_N)$, where

$$\text{Eq. 7) } C_N = \sum_{t=1}^N C_t, \text{ and}$$

$$\text{Eq. 8) } C_t = [(C_1 W_t)$$

Regular Payroll costs from Eq. 1

$$+ C_2 (W_t - W_{t-1})^2$$

Hiring and Layoff costs from Eq. 2

$$+ C_3 (P_t - C_4 W_t)^2 + C_5 P_t - C_6 W_t$$

Overtime costs from Eq. 4

$$+ C_7 (I_t - C_8 - C_9 O_t)^2],$$

Inventory connected costs
from Eq. 5

subject to the restraints,

$$\text{Eq. 9) } I_{t-1} + P_t - O_t = I_t, \quad t = 1, 2, \dots N.$$

The cost function above can be applied to the scheduling decision of a great many factories simply by inserting the appropriate numerical values for the cost parameters: C_1, C_2, \dots, C_9 . When we insert the numerical values that we obtained for the paint factory, Equation 10 is obtained. These numerical values are derived from statistical estimates based on accounting data together with subjective estimates of such intangible costs as delayed shipments to customers. In the interest of simplicity, the influence of the order rate on the optimal inventory level was neglected, i.e., C_9 was set equal to zero.

$$\begin{aligned} \text{Eq. 10) } C_N = \sum_{t=1}^N \bigg\{ & [340W_t] + [64.3 (W_t - W_{t-1})^2] \\ & + [.20(P_t - 5.67W_t)^2 + 51.2P_t - 281.W_t] \\ & + [.0825(I_t - 320)^2] \bigg\} \end{aligned}$$

Where C_N is the total cost for N months expressed in dollars, W_t is the work force for month t expressed in men, P_t is production in gallons (a pseudo-unit to disguise / ^{company cost} data) per month, and I_t is net inventory in gallons.

Since estimates of the cost parameters are subject to many sources of error, it is reassuring that the factory performance proves not to be critically dependent on the accuracy of the cost function. Even if substantial errors are made in estimating the parameters of the cost function, the factory performance measured in cost terms will not suffer seriously.^{1/}

In obtaining the above cost function for the paint factory it should be remembered that the quadratic form of the cost function is an approximation to the "true" cost function. The adequacy of the quadratic approximation can not, however, be judged simply in terms of "goodness of fit." Rather, it must be judged by whether the decisions to which it leads are better than the decisions made by alternative decision methods.

Having translated the decision problem into a precise mathematical problem, we can proceed directly to solve for the best scheduling decisions. Without going into the mathematics involved -- the reader is referred to the technical appendix -- we will now examine the solution that is obtained.

^{1/} An exploratory analysis of the effects of errors in estimating the parameters of a simple quadratic cost function showed that overestimating cost parameters by 100% or underestimating them by 50% -- in both cases estimates were incorrect by a factor of two -- led to decision rules whose cost performance was approximately 11% above the costs which would occur with correct estimates of cost parameters.

IV. The Optimal Decision Rules for the Paint Factory

Once the parameters of the cost function are estimated, the decision rule solution may be obtained by differentiating with respect to each decision variable. We obtain a set of linear equations, and then invert the matrix of these equations to obtain the decision rules. Fortunately, the results of this procedure can be reduced to a formula, requiring only a routine computation. It can be proved mathematically, once and for all, that the decisions yielded by the optimal decision rule are the best possible for the given cost function.

There are two decision rules to be applied at the beginning of each month: one rule sets the aggregate production rate, the other determines the work force. The first rule, shown in Equation 11, incorporates a weighted average of the forecasts of future orders (in this case for a twelve-month period starting with the forthcoming month, t .) Since the forecasts of future orders are averaged, production is smoothed, so that there is an optimal response to the fluctuation of forecasted orders. The weight given to future orders declines rapidly as the forecast extends farther into the future. This occurs because, taking into account the cost of holding inventory, it is not economic to produce currently for shipment in the too remote future. One implication is that there is little point in forecasting orders very far into the future since these orders will have little effect upon optimal current production. For the particular costs of the paint company, the forecasts of orders for the forthcoming and the two successive months are the major determinants of production, as far as orders are concerned.

No information is required about the probability distribution of errors in the forecasts of orders.^{1/} However, the average forecast error should be zero, i.e., the forecasts should be unbiased.

The second term of Equation 11, $(.993 W_{t-1})$, reflects the influence on the scheduled production rate of one of the initial conditions at the time the decision is made -- specifically, the number of workers that were employed at the end of the preceding month. The more workers that are on the payroll at the beginning of the month, the greater should be the production scheduled for the month, since any large decreases in the size of the work force would be costly, as would an unused work force.

The next two terms in the decision rule may be considered together: $(153. - .464 I_{t-1})$. If net inventory at the end of the previous month is large, then the negative term will exceed the positive one, and production will be decreased in order to lower inventory. Similarly, if the initial net inventory is small, the negative term will be small and an increase in production will be called for. Not only does this term determine how the

^{1/} The mathematical analysis indicates that only the expected values of the distributions of orders are relevant to making optimal decisions. The variance and all other higher moments of the distributions have no effect on the decisions under a quadratic criterion. Unbiased forecasts are treated in making decisions exactly as if they were perfect forecasts.

Proofs of this point have been made by C. C. Holt for unrevised forecasts in ONR Research Memorandum No. 3, Carnegie Institute of Technology, "Superposition Decision Rules for Production and Inventory Control," and by H. A. Simon for the general case of revised forecasts in "Dynamic Programming under Uncertainty with a Quadratic Criterion Function" (originally ONR No. 8) which will be published.

optimal production rule responds to any given initial inventory situation, but it has the special significance of indicating how the rule will take account of past forecast errors, since their effect is to raise the net inventory above, or lower it below, the desired level.

$$\text{Eq. 11)} \quad P_t = \left\{ \begin{array}{l} + .463 O_t \\ + .234 O_{t+1} \\ + .111 O_{t+2} \\ + .046 O_{t+3} \\ + .013 O_{t+4} \\ - .002 O_{t+5} \\ - .008 O_{t+6} \\ - .010 O_{t+7} \\ - .009 O_{t+8} \\ - .008 O_{t+9} \\ - .007 O_{t+10} \\ - .005 O_{t+11} \end{array} \right\} + .993 W_{t-1} + 153. - .464 I_{t-1}$$

Eq. 12

$$W_t = .743 W_{t-1} + 2.09 - .010 I_{t-1} + \left\{ \begin{array}{l} + .0101 O_t \\ + .0088 O_{t+1} \\ + .0071 O_{t+2} \\ + .0054 O_{t+3} \\ + .0042 O_{t+4} \\ + .0031 O_{t+5} \\ + .0023 O_{t+6} \\ + .0016 O_{t+7} \\ + .0012 O_{t+8} \\ + .0009 O_{t+9} \\ + .0006 O_{t+10} \\ + .0005 O_{t+11} \end{array} \right\}$$

Where: P_t is the number of units of product that should be produced during the forthcoming month, t .

W_{t-1} is the number of employees in the work force at the beginning of the month (end of the previous month).

I_{t-1} is the number of units of inventory minus the number of units on back order at the beginning of the month.

W_t is the number of employees that will be required for the current month, t . The number of employees that should be hired is therefore $W_t - W_{t-1}$.

O_t is a forecast of number of units of product that will be ordered for shipment during the current month, t .

O_{t+1} is the same for the next month, $t+1$, etc.

The second decision rule, shown in Equation 12, is used to determine the size of the work force. Again, the third term is a weighted average of forecasts of future orders, but in this second rule the weights extend farther into the future before they become negligible in size. Thus the forecasts of orders in the more distant future are relevant in making employment decisions, even though they have little influence on the production decision.

The next term of the employment rule, $.743 W_{t-1}$, indicates that the work force on hand at the beginning of the month will influence employment during the following month, because of the costs associated with changing the work force.

The next two terms in the employment rule, $(2.09 - .010 I_{t-1})$,

incorporate the effect of net inventory on the employment decision. A large net inventory will lead to a decrease in the work force while a small net inventory will tend to require an increase in the work force. Net inventory has a much smaller effect on employment than it has on production. Some general comments can now be made about how these two rules operate in concert.

There is a fairly complex interaction between these two decision rules. The production of one month affects the net inventory position at the end of the month. This in turn influences the employment decision in the second month which then influences the production decision in the third month. Thus there is a continual dynamic interaction between the two decisions.

The influence of net inventory on both the production and employment decisions produces a feedback or self-correcting tendency which eventually returns net inventory to its optimal level regardless of whether or not sales have been forecasted accurately.

The weights that are applied to the sales forecasts and the feedback factors in the two decision rules determine the production and employment responses to fluctuations of orders and thereby indicate how much of these fluctuations should be absorbed by work force fluctuations, overtime fluctuations, and inventory and back order fluctuations in order to minimize costs. The work force responds only to fairly long-term fluctuations in orders, but production responds strongly to the orders in the immediate future and to the inventory position. Thus it appears that short-run fluctuations in orders and the disturbances

that are caused

Page 29.

/by forecast errors are absorbed largely by overtime and undertime fluctuations.

The appearance of negative weights for forecasted future orders in some terms of the production decision rule is surprising. One would expect to prepare for forecasted future orders by increasing production and accumulating inventory. Evidently the response of the rules to a forecast of future orders is to prepare by building up the work force. This in turn gradually leads to increased production.

If the numerical constants in the cost function of the paint factory should change, the numbers in the above decision rules would need to be recomputed in order to obtain new decision rules applicable to the changed circumstances. However, the algebraic forms of the decision rules would remain unchanged.

For procurement or other reasons it may be desirable to know what the production and employment levels are likely to be in subsequent months. Forecasts of future decisions may readily be obtained by means of a set of forecasting rules that are similar to the above decision rules. Of course, when the time comes, the actual decisions may prove to be different from those forecasted.

V. Comparison of Decision Performances by the Factory and the Decision Rules

The decision rules we have described were obtained by finding a mathematical optimum for the decision problem on the basis of specific formal assumptions. In addition, the decision rules have been tested by making a hypothetical application and observing their performance

characteristics. The production and employment decisions that the paint company had made over a six-year period were analyzed in detail. With this knowledge of the decision problems that had confronted the paint factory, the decision rules were applied ex post to simulate the decisions that would have been made if the new decision rules had been used during this period.

Before this hypothetical performance could be calculated, it was necessary to obtain for each point in time a set of forecasts of future orders (in order to calculate the corresponding employment and production decisions for each point in time). Since no explicit forecasts had been recorded by the company, it was impossible for us to operate with the same forecast information that had been available to the factory management at the times when their decisions were made. As a substitute, two different sets of forecasts were computed which, in terms of accuracy, would necessarily bracket the forecasts that were available to the company. The first set of forecasts is the data on orders which were in fact received. Such a Perfect Forecast is of course limited to "forecasting" a known past, and consequently is not of practical usefulness. However, the Perfect Forecast gives a good basis for comparison since, by its use, an upper limit of decision performance is obtained. The second set of forecasts is obtained by assuming that future orders are predicted by a Moving Average of past orders. The total of orders for the coming year is forecasted to be equal to the orders that had been received in the year just past. This forecast

is then converted to a monthly basis by applying a known seasonal adjustment. We now have a basis for a three-way comparison of decision performance: 1) the actual performance of the factory, 2) the performance of the optimal decision rule with Perfect Forecasts whose accuracy cannot be exceeded, and 3) the performance of the optimal decision rule with Moving Average forecasts whose accuracy represents a practical minimum below which there is little excuse for falling.

History of Factory Operations under Alternative Decision Methods

The extreme variability of the orders received by the paint factory is shown in Figure 6. The depressed business conditions of 1949 are clearly reflected in the data. The effects of inventory speculation by distributors and dealers brought on by the Korean War is shown in the high orders of late 1950 and early 1951, and the subsequent rapid decline of orders in the second half of 1951. Hence, the time covered by this study includes a period of extreme order fluctuations as well as periods of more moderate fluctuations. The severity of the fluctuations of orders gives some assurance that the decision rules will be subjected to a test of substantial severity. Although not readily observable by eye, there is a marked seasonal pattern in the receipt of orders.

An examination of Figure 7 shows that the production fluctuations of the factory are considerably sharper on a month-to-month basis than those called for by the decision rule with either Moving Average or

FIGURE 6





FIGURE 7




Perfect Forecasts.^{1/} With a Perfect Forecast the decision rule avoids, almost completely, sharp month-to-month fluctuations in production, but responds to fluctuations of orders that have a duration of several months.

The decisions scheduling the size of the work force are shown in Figure 8. Again, the decision rule makes smoother changes and avoids sharp month-to-month fluctuations in work force. The fluctuations in work force with the Perfect Forecast, while substantial in size, are actually occasioned by the severity of order fluctuations and the desire to avoid costly accumulations of inventory and back orders. The additional work force fluctuations^{that} are observed under the Moving Average forecast are entirely attributable to forecast errors. For example, an erroneous forecast of high sales leads the decision rule to build up the work force. The combination of low sales and large work force causes an accumulation of inventories which in turn necessitates a reduction of the work force in order to lower inventory to the optimal level. The differences which are shown in Figure 8 between the fluctuations of the work force under the Perfect Forecast and the Moving Average forecast when the same decision rule is used in both cases illustrate the importance of accurate forecasts to the stability of employment.

As would be expected, the Perfect Forecast foresaw the increased

^{1/} For the factory no adjustment was made for the fact that the number of working days varies somewhat from month to month. This accounts for part of the production variability.

FIGURE 8



"Korean" orders and increased the size of the work force sharply in 1950. Using the Moving Average forecasts, the decision rule increased the work force about six months later. While the factory actually started its employment buildup as early as the decision rule did using Perfect Forecasts, its rate of buildup was considerably lower; consequently its peak of employment occurred in late 1951 at the time when, as it happened, orders declined sharply. Evidently the decision rule when using the Moving Average forecast worked tolerably well even under such severe circumstances as the outbreak of war.

Overtime hours are plotted in Figure 9 to show the comparisons in performance between the factory and the decision rule. The inadequacies of the Moving Average forecast appear clearly in 1950, when the sudden war-induced increase in orders, which, of course, were not foreseen by the backward-looking forecast, led to a large amount of overtime. ^P Performance in the control of inventory is shown in Figures 10 and 11, which show separately the two components of net inventory; actual physical inventory and back orders. The decision rule operating with the Perfect Forecast displays in Figure 10 the ability to hold inventories quite close to the lowest cost level. Deviations from this optimal level do occur, but they are not of large amplitude. In contrast, the decision rule operating with the Moving Average forecast allows inventories to fall substantially during the sudden increase in "Korean" sales, and later, when orders decline, inventory rises sharply. However, inventory recovered from its low point much earlier

FIGURE 9





FIGURE 10




with the Moving Average forecast than the factory actually did. In the winter of 1951-52 when orders declined sharply, the decision rule using the Moving Average forecast was able to bring down the resulting excess inventories about as quickly as this was in fact achieved by the factory.

The penalty that accompanies low inventory appears clearly in the plot of back orders in Figure 11. With the Moving Average forecast, back orders rose sharply during the Korean spurt of demand, but these back orders were liquidated by the end of 1950. For the actual factory performance, back orders did not return to their normal level until the second half of 1951. When high orders are speculative in nature as was the case during this period, it is difficult to judge how much weight should be attached to the poor service to customers evidenced by large back orders. The decision rule "took" these back orders seriously and responded accordingly.

Cost Comparisons under Alternative Decision Methods

One test of a decision-making process is its performance in terms of the criteria that serve as the basis for the decisions. To the extent that the minimization of the costs which occur in the cost function constitutes the goal of the production executives of the paint factory, the comparison between the cost performances of the factory and of the decision rule calculated on this basis is significant. However, the production executives have been concerned during this six-year period with the accomplishment of other goals in addition to the minimization

FIGURE 11



of the particular set of costs with which the statistical decision analysis is concerned; pursuit of these other goals would undoubtedly raise these costs. Hence performance comparisons based exclusively on the costs that are included in the cost function do not tell quite the whole story.

Because the reconstruction of a quantitative history of factory operations for six years constitutes in itself a substantial research job involving in this case the allocation of costs between paint and non-paint, the indirect calculation of certain information that had never been recorded, and the estimation of nonaccounting costs, the figures that have been obtained must be presented with a certain tentativeness. Similarly the estimates of what the costs would have been, if things had been done differently, are peculiarly subject to limitations in accuracy.

In spite of their limitations, the cost comparisons to be presented are, in the opinion of the authors, highly significant.

To evaluate the cost performance of the decision rules, including the adequacy of the fit of the quadratic cost function, we used, so far as possible, the cost structure that originally had been estimated from the factory accounting and other data.

A cost comparison is shown in Figure 12, for 1949-53, the longest period in which cost figures are available for a complete three-way comparison. The year 1954 could not be included, because, at the time of writing, the authors could not produce the Perfect Forecast of 1955 orders which would be required.

COST COMPARISONS FOR 1949~1953

FIGURE 12

COSTS (THOUSANDS OF DOLLARS)	COMPANY PERFORMANCE	DECISION RULE	
		MOVING AVG. FORECAST	PERFECT FORECAST
REGULAR PAYROLL	\$ 1940	\$ 1834	\$ 1888
OVERTIME	196	296	167
INVENTORY	361	451	454
BACK ORDERS	1566	616	400
HIRING & LAYOFFS	22	25	20
TOTAL { \$		3222	2929
COST %		110%	100%

The decision rule with Perfect Forecasts had lower costs than with the Moving Average Forecasts by \$59,000 per year on the average. Since the identical rule is being used with both sets of forecasts, this difference in cost performance is entirely attributable to better forecasting. The decision rule when operating with the obviously modest forecasting ability of the Moving Average gave a cost saving compared to the factory performance of \$173,000 per year on the average. The limitations of this comparison which were mentioned above should be noted.

It is striking that the cost saving attributable to the decision rule is greater than the cost saving attributable to errorless forecasting. Perhaps forecasting future orders accurately isn't as important as has commonly been thought by production people. Judging by this particular factory and period, making optimal use of crude forecasts is more important than perfect forecasting.

Since scheduling production and employment in a period of recession and war is so difficult a problem due to the large and unpredictable fluctuations of orders, it is understandable that the potential savings through improved decision techniques should be large. However, the Korean war period may by some be considered unrepresentative of attainable cost savings in hoped-for times of "normalcy." The Korean period also presents difficult problems in estimating an appropriate penalty cost for back orders. Cost comparisons for shorter periods that exclude the war years should be more representative of "normal"

times. However, in posing an easier scheduling problem, these years, of course, offer smaller opportunities for improved performance.

If we drop out the Korean year, 1950, and compare the Perfect Forecast cost performance with that of the Moving Average for the years 1949, 1951, 1952 and 1953, we find that the imperfect forecasting raises costs 5%, or \$28,000 per year on the average. While a 5% savings is small in percentage terms, it should be remembered that this is 5% of an amount that is the total of several large costs including the payroll. How much of this saving can actually be achieved by substituting more refined forecasting methods for the moving average is as yet unknown. Obviously a perfect forecasting method is unattainable. Presumably the expenditure of some thousands of dollars for improved forecasts would more than pay for itself in decreased production costs even for this small one-hundred-man paint factory.

Although the expected size of forecast errors for a particular forecasting method does not affect optimal decisions based on its forecasts, the cost performance certainly is affected by the size of the forecast errors. Since the cost function is quadratic, the costs of forecast errors rise roughly with their square. Hence it is desirable

to find a forecasting method that does not often make large errors^{1/} — small errors can be forgiven because their cost penalty is low.

The plot of actual factory inventory in Figure 10 shows that the factory in its control of inventories acted as if the cost of back orders relative to the cost of holding inventories had increased during this six-year period. The cost structure that we estimated is more nearly in line with the implicit back order and inventory costs of the later three-year period. Consequently, cost comparisons from the later period may be more significant than cost comparisons covering the whole six years.

The objection may occur to the reader that our estimates of the factory cost structure may be in error which would mean that the factory performance is being judged by an erroneous criterion. Such errors clearly are possible, but it should be remembered that the decision rule is designed to minimize a given cost function. If the cost parameters were changed, the costing of factory performance would be

1/ The paper by C. C. Holt mentioned in Footnote 1 on Page 25 presents a partial analysis of the cost of forecast errors.

In an analysis of production scheduling (not considering employment) on the basis of an unrevised forecast of orders that will occur in one future period, the expected cost of forecast errors was found to be proportional to the variance of the distribution of forecast errors. The constant of proportionality is the square of one of the decision rule weights (that are applied to order forecasts) corresponding to the time spanned by the forecast. Not only do forecasts of future orders have less influence on decisions when they are more remote in the future, but the cost implications of forecast errors attenuate even faster (as the square of the weight) as the forecast reaches farther into the future.

different, but also a new decision rule would be calculated whose decision behavior would be different. Consequently if changes were made in the cost structure that would reduce the estimated cost of the factory performance, the relevant comparison would then be with the cost performance of a decision rule changed to be optimal under the new cost function.

To compare the cost performances of the factory and the decision rule with Moving Average forecasts we chose the period 1952-54 — the latter year is available for this comparison since the Moving Average forecasts require no unattainable data on 1955 orders. As shown in Figure 13, the actual factory cost performance exceeded that of the decision rule with Moving Average forecasts by 8.5%, or \$51,000 per year on the average. Economics were achieved by the decision rule as follows: The overtime costs under the decision rule were higher, but the regular payroll costs were enough lower to make a net saving. The inventory holding costs were higher under the decision rule, but the back order penalty costs were enough lower to make a net saving. The hiring and layoff costs were lower under the decision rule. It appears that the cost savings during this period of "normal" paint sales were attained by the decision rule through a combination of several different kinds of cost savings and not through a single simple improvement that might be "hit upon" by casual judgmental analysis.

COST COMPARISONS FOR 1952~1954

FIGURE 13

COSTS (THOUSANDS OF DOLLARS)	COMPANY PERFORMANCE	DECISION RULE MOVING AV. FORECAST
REGULAR PAYROLL	\$ 1256	\$ 1149
OVERTIME	82	95
INVENTORY	273	298
BACK ORDERS	326	246
HIRING & LAYOFFS	16	12
TOTAL	\$ 1953	\$ 1800
COST	108.5%	100%

VI. Conclusion

On the basis of the foregoing comparisons between the actual decision performance of an operating factory and the hypothetical performances of the decision rule, the following conclusion seems justified. If the optimal linear decision rule which is introduced in this paper were to be applied using forecasts that are practically obtainable, it would render a performance that would be a considerable improvement in cost terms over that obtainable by the traditional judgmental methods that have been used by the factory. Furthermore, this improved performance probably could be obtained with a smaller expenditure of executive time and effort than now goes into such decisions.

It would be rash of the authors to generalize these conclusions to industry generally, but on the basis of their knowledge of the decision techniques that are now in general use, it is their opinion that the decision performance of the paint factory is not atypical, and that the optimal linear decision rule which is presented in this paper would, in a great many industrial situations, enable production executives to achieve a substantial improvement in their production and employment scheduling.

Even though a production executive may be aided by adopting this new decision technique, there is still critical need for his judgment, both in the estimation of the original cost function, especially the intangible components of it, and in the application of the decision rule when factors become important that are not explicitly included in

the statistical decision analysis. By relieving the executive of the recurring need to consider and analyze the complex interacting cost factors that are taken into account by the decision rule (i.e., regular payroll, overtime, hiring, layoff, inventory, back order and machine setup costs), the executive will be left with more time free to devote attention to important nonroutine special factors and unusual situations.

Even though it is possible to prove mathematically that, where the cost function is quadratic, the linear decision rule here presented cannot be surpassed on its average cost performance, and even though a simulated application based on historical data gives highly encouraging results, there is still need for further tests under actual operating conditions before the new method can fully prove its usefulness. For the last year, the paint factory has been carrying out an application of the rule (and an earlier version of it) to the actual scheduling of its production in order to test the rule under operating conditions with available forecasts. The results have been gratifying. The average inventory and average back orders have both decreased, and this was accomplished with smaller fluctuations in the aggregate production of those paints that were included in the experiment.^{1/}

^{1/} A report on this operating test is in preparation for publication in the Naval Research Logistics Quarterly. This test also involves the application of decision rules for scheduling the production of individual products as well as the rules that have been the subject of this paper for scheduling aggregate production.

Since it is anticipated that the decision rule discussed in this paper will be applicable in its present form to a good many other factories, a technical appendix is attached which shows the derivation of the rule and the formulas for obtaining the final decision rules for scheduling production and employment.

Other factories will need to include different types of cost terms in their cost functions and this will preclude using the rule in the form in which it has been presented. However, the general method of obtaining a linear decision rule from a quadratic criterion function can be extended by the application of straightforward mathematics to a wide variety of other production and inventory control decision problems, and, of equal significance, problems from entirely different fields.^{1/}

^{1/} A research memorandum presenting general methods for deriving optimal linear decision rules for quadratic criterion functions is in preparation.